Distance Bounds for High Dimensional Consistent Digital Rays and 2-D Partially-Consistent Digital Rays

Man-Kwun Chiu 🄳
Institut für Informatik, Freie Universität Berlin, Germany
chiumk@zedat.fu-berlin.de

Matias Korman
Department of Computer Science, Tufts University, Medford, MA, USA
matias.korman@tufts.edu

Martin Suderland 🄳
Faculty of Informatics, Università della Svizzera italiana, Lugano, Switzerland
martin.suderland@usi.ch

Takeshi Tokuyama
Kwansei Gakuin University, Sanda, Japan
tokuyama@kwansei.ac.jp

Abstract

We consider the problem of digitalizing Euclidean segments. Specifically, we look for a constructive method to connect any two points in $\mathbb{Z}^d$. The construction must be *consistent* (that is, satisfy the natural extension of the Euclidean axioms) while resembling them as much as possible. Previous work has shown asymptotically tight results in two dimensions with $\Theta(\log N)$ error, where resemblance between segments is measured with the Hausdorff distance, and N is the L_1 distance between the two points. This construction was considered tight because of a $\Omega(\log N)$ lower bound that applies to any consistent construction in $\mathbb{Z}^2$.

In this paper we observe that the lower bound does not directly extend to higher dimensions. We give an alternative argument showing that any consistent construction in d dimensions must have $\Omega(\log^{1/(d-1)} N)$ error. We tie the error of a consistent construction in high dimensions to the error of similar *weak* constructions in two dimensions (constructions for which some points need not satisfy all the axioms). This not only opens the possibility for having constructions with $o(\log N)$ error in high dimensions, but also opens up an interesting line of research in the tradeoff between the number of axiom violations and the error of the construction. In order to show our lower bound, we also consider a colored variation of the concept of discrepancy of a set of points that we find of independent interest.

1 Introduction

Euclidean line segments are one of the most fundamental objects of geometry. Although often loosely referred to as *the shortest path connecting the endpoints*, segments have a clear and unique axiomatic definition out of which many interesting properties follow. For example, it is well-known that the intersection of two segments is always a segment (that could possibly degenerate to a point or even become empty). The definition of other mathematical concepts heavily depends on the definition of segments (such as convex regions).

The definition of segment works very well in a Euclidean or similar spaces with infinite precision. Digital representation (such as pixels in a screen) introduces imprecision. The most common approach used in practice is to somehow *round* the Euclidean segment into the digital space. The digital segments will look very similar to the Euclidean counterparts (that is, the *error* is very small). However, we cannot guarantee the useful properties and concepts that follow from the axiomatic definition of Euclidean segment (see Figure 1).

In the aspect of the consistency of digital segments, we look for a deterministic method to construct digital segments in a way that (*i*) the analogous of Euclidean axioms are satisfied and (*ii*) the digital segments resemble the Euclidean ones as much as possible.
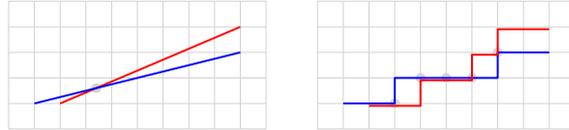
Figure 1 Left: Two Euclidean line segments that intersect in a point. Right: Rounding produces polylines that intersect in three disconnected components.

Preliminaries

Our aim is to construct a digital path $dig(p,q)$ for any two points $p, q \in \mathbb{Z}^d$. Ideally, we want dig to be defined for any pairs of points in $\mathbb{Z}^d$ (full list of requirements is described below), but sometimes we consider the case in which dig is only defined for a subset of $\mathbb{Z}^d \times \mathbb{Z}^d$.

▶ **Definition 1.** *For any $S \subseteq \mathbb{Z}^d \times \mathbb{Z}^d$, let $DS(S)$ be a set of digital segments such that $dig(p,q) \in DS(S)$ for all $(p,q) \in S$. We say that $DS(S)$ forms a* partial set of consistent digital segments *on S (partial CDS for short) if for every pair $(p,q) \in S$ it satisfies the following five axioms:*
(S1) *Grid path property: $dig(p,q)$ is a path between p and q under the $2d$-neighbor topology[1].*
(S2) *Symmetry property: if $(q,p) \in S$, $dig(p,q) = dig(q,p)$.*
(S3) *Subsegment property: for any $r \in dig(p,q)$, $dig(p,r) \in DS(S)$ and $dig(p,r) \subseteq dig(p,q)$.*
(S4) *Prolongation property: $\exists\, r \in \mathbb{Z}^d$ such that $dig(p,r) \in DS(S)$ and $dig(p,q) \subset dig(p,r)$.*
(S5) *Monotonicity property: for all $i \leq d$ such that $p_i = q_i$, it holds that every point $r \in dig(p,q)$ satisfies $r_i = p_i = q_i$.*

These axioms give nice properties of digital segments analogous to Euclidean line segments. For example, (S1) and (S3) imply that the intersection of two digital segments is another segment (that could degenerate to a single point or an empty set). (S5) implies that the intersection of a segment with an axis-aligned halfspace is a segment.

---

[1] The $2d$-neighbor topology is the natural one that connects to your predecessor and successor in each dimension. Formally speaking, two points are connected if and only if their $\|\cdot\|_1$ distance is exactly one.

A partial CDS for $S = \mathbb{Z}^d \times \mathbb{Z}^d$ is called a set of *consistent digital segments* (CDS for short). Although our final goal is to have such a construction that works for the case in which $S = \mathbb{Z}^d \times \mathbb{Z}^d$, in this paper we consider subsets of the form $S = \{o\} \times \mathbb{Z}^d$ (where o is the origin or any fixed point in $\mathbb{Z}^d$). We say that a partial CDS on such a set is a *consistent digital ray* system (CDR for short), as it contains all segments (or rays) from o to $\mathbb{Z}^d$.

Another property that we want from partial CDS is that they visually resemble the Euclidean segments. The resemblance between the digital segment $dig(p,q)$ and the Euclidean counterpart $\overline{pq}$ is measured using the *Hausdorff* distance. The Hausdorff distance $H(A,B)$ of two objects A and B is defined by $H(A,B) = \max\{h(A,B), h(B,A)\}$, where $h(A,B) = \max_{a \in A} \min_{b \in B} \delta(a,b)$, and $\delta(a,b)$ is the standard $|| \cdot ||_\infty$ L-infinity norm.

Thus, the resemblance of a partial CDS on S is defined as $\max_{(p,q) \in S} H(dig(p,q), \overline{pq})$ (that is, the biggest error created between a digital segment and its Euclidean counterpart). This value is simply referred to as the *error* of the partial CDS construction. We are interested to see how the error grows as we enlarge our focus of interest. Thus, we limit the domain to the case in which both points are in the L_1 ball of radius N centered at the origin (i.e. $\mathcal{G}_N = \mathbb{Z}^d \cap B_1(o, N)$). Rather than looking for the exact function, we are interested in the asymptotic behavior of the error as a function of N. For simplicity, we will actually restrict ourselves to the positive orthant $\mathcal{G}_N^+ = \mathcal{G}_N \cap_i H_i$, where $H_i = \{p \in \mathbb{Z}^d : p_i \geq 0\}$ and p_i is the i-th coordinate of p (the results extend to other orthants by symmetry).

Previous Work

Research on the digital representation of line segments has been an active area of research for over half a century [10]. Many different approaches have been considered. Most common techniques look for methods that implicitly encode the properties we desire. For example, a popular approach is to consider a *dynamic* method to digitize line segments. In this setting, the way we transform a Euclidean segment into a digital one will depend on which other segments are present (and their specific coordinates). It is known that a grid of exponential size is needed if we want to preserve the combinatorial types [9]. Another workaround is known as *snap rounding* that represents line segments by polygonal chains: Each segment is carefully rounded to avoid inconsistencies. Note that both of these ideas implicitly keep the error small while making sure that the intersection of two digital segments is a connected component. Although they work well in practice, they have the drawback that they cannot be used to define objects that are based on digital segments (such as digital starshapes or convex region).

The first paper to explicitly look for an axiomatic approach was in 1987 by Luby [11]: in his work he introduced the concept of CDS (under the name of *smooth geometries*) and gave a method to construct CDSs in $\mathbb{Z}^2$ based on a characterization of CDRs in $\mathbb{Z}^2$: any CDR can be uniquely identified by four total orders of the integers (and vice versa). By choosing a proper total order and using it for all points of $\mathbb{Z}^2$ we obtain a CDR with $O(\log N)$ error. Håstad[2] gave a matching lower bound for any such construction. The lower bound is based on discrepancy theory [12]: any CDR is mapped to a sequence of real numbers in $[0, 1)$ in a way that the error of the CDR is proportional to the discrepancy of the sequence (intuitively speaking, a measure on how well shuffled the numbers are).

These results were rediscovered by Chun et al. [8] and Christ et al. [6]. They renewed interest in the topic and sparked other related research: Chowdhury and Gibson [4] gave necessary and sufficient conditions for a collection of CDRs to form a CDS. In a companion

---

[2] The lower bound was published by Luby, but credit given to Håstad (see Theorem 19 of [11]).

paper, the same authors [5] afterwards provided an alternative characterization together with a constructive algorithm; specifically, they gave an algorithm that, given a collection of segments in an $N \times N$ grid that satisfies the five axioms, computes a CDS that contains those segments. The algorithm runs in polynomial time of N.

Unfortunately, most of these results only work on the digital plane. Out of the previously mentioned results, only the CDR construction of Chun et al. [8] extends to three and higher dimensions. The construction has $O(\log N)$ error regardless of the dimension. Chun et al. [8] also considered the case in which the monotonicity property (S5) is not preserved. They showed that if we remove (S5), we can obtain a CDR with $O(1)$ error in any dimension. Although the error is small, the resulting segments are far from what we would consider similar to the Euclidean segments (because they loop around many times). Recently, Chiu and Korman [2] showed that the problem in higher dimensions behaves very differently from the two dimensional case. Specifically, they studied how to extend the CDS construction of Christ et al. [6] and showed that it is very limiting in three (and higher) dimensions. We can use their method to get arbitrarily many CDRs (with $\Omega(\log N)$ error) and sometimes we can get a CDS. However, whenever the construction yields a CDS, it will have $\Omega(N)$ error.

Our interest in higher dimensions comes motivated by an application in *image segmentation*. Image segmentation is the act of separating an object from its background in an image (that is, determining which pixels are part of the background and which ones not). Chun et al. [8] showed how to combine their CDR construction with the framework of Asano et al. [1] to segment two dimensional images. This idea has been extended to consider other shapes (see [7] for a detailed list), but always two dimensional. The hope is that a high dimensional CDR with low error will produce more accurate segmentation algorithms. Although traditional images taken with a camera are two dimensional, images from a medical equipment such as those taken with an MRI machine can have three or even higher dimensions (say, when we want to track changes of a particular object along time).

Results and paper organization

When approximating some geometric object, it often happens that higher dimensions create a larger error than in a lower dimension setting. Since the high dimensional setting contains a two dimensional subspace, it is common for lower bounds to extend to higher dimensions. However, this is not true for the case of CDRs: although a three dimensional CDR contains two dimensional subspaces, those subspaces need not exactly be CDRs (and thus the $\Omega(\log N)$ lower bound does not directly hold). In this paper, we further explain the reason and investigate the lower bound for the higher dimensional case.

The main reason why a subspace is not a CDR is because of the prolongation property (S4): we require that every segment is extendable, but has no constraints on the dimension in which it does so. In particular, a subspace of a high dimensional CDR need not be a CDR (see an example in Figure 2). Subspaces of CDRs are what we call *weak CDR*: it is a construction that almost always behaves like a CDR but some vertices may not satisfy the prolongation property (S4). Each vertex that does not extend is called an *inner leaf*. In this paper we study weak CDRs in two dimensions and the implications that they have for (proper) CDRs in higher dimensions.

The new found properties of weak CDRs allow us to extend the two-dimensional lower bound to higher dimensions. Håstad's bound was based on a mapping from a (two-dimensional) CDR into a pointset in $[0,1) \subset \mathbb{R}$ and tied the error of the CDR to the discrepancy of the transformed pointset. Our lower bound uses an additional intermediate step: from any CDR we consider the weak CDR it generates in the $x_1 x_2$-plane. We then
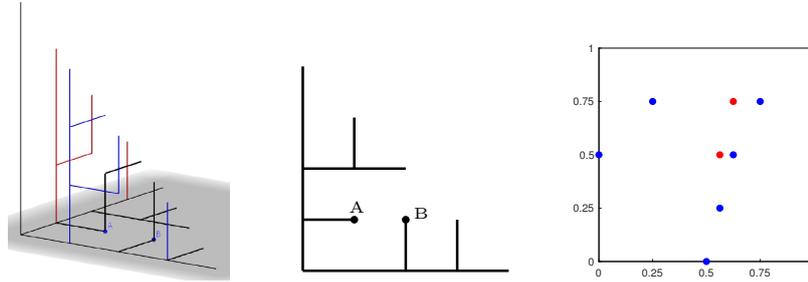
Figure 2 (left) A drawing of a CDR in $\mathcal{G}_N^+ \subset \mathbb{Z}^3$ for $N = 4$. Notice that the CDR is a tree whose leaves are at the plane $x + y + z = N$. (middle) A cross section on the xy-plane of the same CDR. Observe that vertices A and B do not extend within the xy-plane. Thus, the subspace is a weak CDR (rather than a proper CDR). (right) A map of the weak CDR into a two-colored pointset. Regions with many blue points and few red correspond to portions of the CDR with high error.

map this weak CDR into a set of points in the unit square and *then* use discrepancy theory to obtain a lower bound for the weak CDR and eventually to the high dimensional CDR. Overall, we show a very strong link between the three spaces (CDR in high dimensions, weak CDR in the $x_1 x_2$-plane and set of points created by our mapping). Along the paper we will analyze properties of each of the spaces, and see what implications it has for the other two. Specifically, we show the following:

(i) Because we now need to account for more general constructions (weak CDRs instead of proper CDRs), the mapping needs to be changed. Instead of creating points in the $[0, 1)$ interval, in Section 2 we map into a two-colored pointset in $[0, 1) \times [0, 1)$.

(ii) Similar to the two dimensional case, we can tie the error of the weak CDR to the discrepancy of the mapped pointset. First, we extend the discrepancy results [12] to our exact setting. Let R and B be a set of red and blue points in the unit square, respectively. Let $m = |B| - |R|$ and assume $m > 0$. For any set P of points in the unit square and $x, y \in [0, 1]$ let $P[x, y]$ be the number of points in $P \cap [0, x] \times [0, y]$. For any two real numbers $0 \leq x, y \leq 1$ we define the discrepancy of R and B at (x, y) as $D_{R,B}(x, y) = mxy - (B[x, y] - R[x, y])$. The *discrepancy* of R and B is simply defined as $D_{R,B}^* = \max_{(x,y) \in [0,1]^2} |D_{R,B}(x, y)|$ (i.e., the highest discrepancy we can achieve among all possible rectangles). The discrepancy $D_{R,B}^*$ of a two-colored pointset is high if and only if there is an axis-aligned rectangle with the origin as corner in which the difference of the cardinalities is far from the expected difference.

▶ **Theorem 2** (Two colors discrepancy). *For any set R and B of points such that $|B| > |R|$ it holds that $D_{R,B}^* = \Omega\left(\frac{(|B|-|R|) \cdot \log(|B|+|R|)}{|B|+|R|} \right)$.*

The proof is given in Section 3.

(iii) With this new discrepancy result we obtain a trade-off between the error of any weak CDR and the number of inner leaves (i.e., vertices that do not satisfy (S4)). When the weak CDR has zero inner leaves (and thus is a proper CDR) our bound matches the lower bound of Håstad. As the number of inner leaves increases, the lower bound decreases. In Section 4 we prove the following relationship.

▶ **Theorem 3.** *For any $N \in \mathbb{N}$, any weak CDR defined on $\mathcal{G}_N^+ \subset \mathbb{Z}^2$ with κ_2 inner leaves between lines $x + y = \lceil N/2 \rceil$ and $x + y = N$ has $\Omega(\frac{N \log N}{N + \kappa_2})$ error.*

(iv) We then apply Theorem 3 to obtain a lower bound for CDRs in d dimensions: intuitively speaking, if the 2-D subspace has few inner leaves (say, $o(N \log N)$), then it will have $\omega(1)$ error. On the other hand, a weak CDR with many inner leaves in the 2-D

subspace will cause too many points to extend to one of the remaining dimensions, and create large error as well. This gives a lower bound of $\Omega(\log^{1/(d-1)} N)$ for any CDR construction in d dimensions (see Section 5):

▶ **Theorem 4.** *Any CDR in $\mathbb{Z}^d$ has $\Omega(\log^{1/(d-1)} N)$ error.*

Although we believe our analysis to be loose (especially in Theorem 4), we are not certain that the existing CDR constructions with $O(\log N)$ error are tight either. In the full version of the paper in [3], we explore the possibility of having a CDR in high dimensions with $o(\log N)$ error (rather than directly looking at CDRs in high dimensions, we see what properties it would imply in the other two subspaces). Although we cannot explicitly find a construction with $o(\log N)$ error, we provide interesting insight on how further research can solve this question. In particular, we give a weak CDR construction with $5/2$ error and $\Theta(N^2)$ inner leaves. In order to further reduce the number of inner leaves in weak CDRs with constant error we instead look at how to create a two-colored pointset with constant discrepancy. We show that it is not possible to have $o(N^2)$ red points in some pattern of the pointset with constant discrepancy, which gives us a condition on any weak CDR with $o(N^2)$ inner leaves.

Further discussion on the implication of these results is given in Section 6.

2 Mapping a weak CDR into a pointset

We start by showing how to transform a weak CDR in two dimensions into a two-colored pointset in $[0, 1)^2$. Given any weak CDR, its restriction to $\mathcal{G}_N^+$ forms a spanning tree T of $\mathcal{G}_N^+$ because of axioms (S1) and (S3). Although the tree is undirected, we see it as a directed graph (rooted tree) whose edges are oriented away from the origin (root). Then, (S5) implies that the parent of each vertex (x, y) (except the root) is either $(x - 1, y)$ or $(x, y - 1)$. For any edge $e = uv$ of T, where u is the parent of v, we define $T(e)$ as the subtree of T that is rooted at the child node v of e. We slightly abuse the notation and use $T(v)$ to denote the subtree that is emanating from v towards the leaves (that is, $T(v) = T(e)$).

For any $n \leq N$ let L_n be the points of $\mathcal{G}_N^+$ whose sum of coordinates is n (i.e., $L_n = \{(x, y) \in \mathcal{G}_N^+ : x + y = n\}$). We follow the usual terminology that we call a vertex of degree one a *leaf*. We further consider two subcategories: we say that a leaf v of T is an *inner leaf* if it is not in L_N. All the vertices in L_N are called *boundary leaves*. Note that, by properties of CDR, all vertices of L_N are proper leaves (since any children should be in L_{N+1}, which is outside $\mathcal{G}_N^+$). Further note that in a proper CDR there will be no inner leaves. A vertex v of T is a *split vertex* if it has degree three or it is the origin. Let $\mathcal{S}$ be the set of split vertices and $\mathcal{D}$ the set of inner leaves.

2.1 Auxiliary function

Before giving the transformation from a tree to a point set we first define an auxiliary function $M : \mathcal{G}_N^+ \to [0, 1]$. For any $p \in L_N$ we set $M(p) = \frac{p_x}{p_x + p_y}$. For any subtree $T'(v)$ of T we define two more functions inductively for $v \in L_n$ from $n = N$ to 0 as follows:

$$\max(T'(v)) = \max_{p \in T'(v) \cap (\mathcal{D} \cup L_N)} M(p) \quad \text{and} \quad \min(T'(v)) = \min_{p \in T'(v) \cap (\mathcal{D} \cup L_N)} M(p),$$

where $M(p)$ for $p \in \mathcal{D}$ is defined in the next paragraph.

For any inner leaf $\ell \in \mathcal{D}$, we know that the edges $e_1 = (\ell_x - 1, \ell_y + 1)(\ell_x, \ell_y + 1)$ and $e_2 = (\ell_x + 1, \ell_y - 1)(\ell_x + 1, \ell_y)$ must be present in T. Thus, we define $M(\ell)$ as $M(\ell) = \frac{\max(T(e_1)) + \min(T(e_2))}{2}$. Intuitively speaking, we look at the leaves above and to the right of ℓ, and assign a value that is in between the two of them (see Figure 3, left). The following statement shows that these values are sorted along L_n.

▶ **Lemma 5.** *Let $T(u), T(v) \subset T$ be two subtrees of T rooted at the vertices $u, v \in L_n$ (respectively) for some $n \leq N$ such that $u_x < v_x$. Then, it holds that $\max(T(u)) < \min(T(v))$.*

Proof. We prove this statement by induction on n from N to 1. If both $u, v \in L_N$ then both $T(u)$ and $T(v)$ consist of a single vertex and the proof trivially follows. Now, assume that the statement is true for any two vertices $u', v' \in L_i$ for $i > n$. We need to show that the statement holds for any two vertices $u, v \in L_n$ such that $u_x < v_x$.

First observe that if we have two descendants u' and v' from u and v respectively such that $u', v' \in L_{n'}$ for some $n' > n$, then it holds that $u'_x < v'_x$. Indeed, this follows from the fact that when we embed T in the natural way with edges drawn as straight segments, the result is a tree with no crossings. Thus, if $v'_x < u'_x$ happened for some descendants, then the two paths in T from u to u' and from v to v' would either cross or form a cycle. Any of those two situations would contradict with the fact that T is a weak CDR.

Back to our original proof, consider the case in which neither u nor v are inner leaves. By the above argument we have that the x-coordinate of any child $u' \in L_{n+1}$ of u must be smaller than any child $v' \in L_{n+1}$ of v. By induction, this implies that $\max(T(u')) < \min(T(v'))$ and thus $\max(T(u)) < \min(T(v))$.

The cases in which u or v are inner leaves are similar: if u is an inner leaf, we have $\max(T(u)) = M(u) = \frac{\max(T(u_1)) + \min(T(u_2))}{2}$, where $u_1 = (u_x, u_y + 1) \in L_{n+1}$ and $u_2 = (u_x + 1, u_y) \in L_{n+1}$. By induction on u_1 and u_2 we have $\max(T(u_1)) < \min(T(u_2))$ and $\max(T(u)) < \min(T(u_2))$, thus we need to compare $\min(T(u_2))$ with any children of v. If v is also an inner leaf, we can do a similar argument and have that $\max(T(v_1)) < \min(T(v))$ where $v_1 = (v_x, v_y + 1)$.

In general, given u, let $u' \in L_{n+1}$ be the child of u with the largest x-coordinate (or $u' = u_2$ if u is an inner leaf). Similarly, we define v' as the child of v with the smallest x-coordinate (or $v' = v_1$ if v is an inner leaf). Again, by planarity of the natural embedding, we have that $u'_x \leq v'_x$ if at least one of u, v is an inner leaf. In either case, we can use induction and get that $\max(T(u')) \leq \min(T(v'))$ which implies $\max(T(u)) < \max(T(u')) \leq \min(T(v')) \leq \min(T(v))$ (if u is an inner leaf) or $\max(T(u)) \leq \max(T(u')) \leq \min(T(v')) < \min(T(v))$ (if v is an inner leaf) completing the proof. ◀

For any subtree T' of T, its *depth* is the longest possible length of a path from its root to any of its leaves. Any split vertex $s \in \mathcal{S}$ has two branching edges e_1 and e_2, each defining a subtree. The subtree of higher depth is the *preferred* subtree of s (in case of tie, we choose the tree emanating from $(s_x + 1, s_y)$). For any point $p \in \mathcal{G}_N^+$ we define a walk from p to some leaf of T. If $p \in L_n$ has degree two, we follow the single edge to L_{n+1}. If $p \in \mathcal{S}$, we follow the edge to the preferred subtree. This process is continued until we reach a leaf $\gamma(p)$.

With this virtual walk we can define the function M to all points $p \in \mathcal{G}_N^+$ (not only leaves) of the domain as follows. If p is neither a split nor a leaf, we define $M(p)$ as $M(p) = M(\gamma(p))$. For a split vertex s, let s' be the child of s that is not on the preferred subtree of s. Then, we define $M(s)$ as $M(s) = M(\gamma(s'))$.

Intuitively speaking, from any vertex we always follow its only edge away from the root (if it has degree 2) or the preferred edge (if it has degree 3) until we reach a leaf. The only exception is if we start on a split vertex, in which case we do not follow the preferred edge at the first step. This exception is needed to make sure that the end points of the walk starting from split vertices are distinct.

▶ **Lemma 6.** *For any split vertex $s \in \mathcal{S}$, there exists a unique leaf $\ell \in \mathcal{D} \cup L_N$ such that $M(s) = M(\ell)$. And for any leaf $\ell \in \mathcal{D} \cup L_N \setminus \{(N, 0)\}$, there exists a unique split vertex $s \in \mathcal{S}$ such that $M(s) = M(\ell)$.*
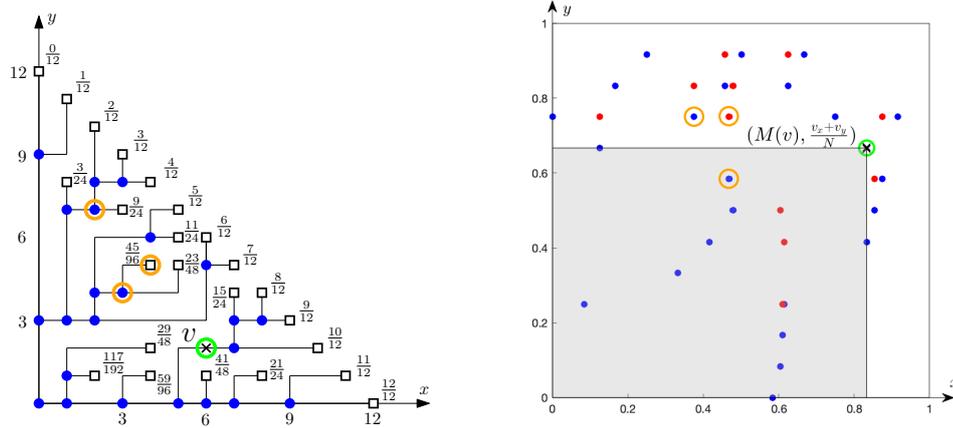
Figure 3 (left) A tree of a weak CDR and the value of the auxiliary function M applied to all leaves of the tree. (right) The tree transformed into blue and red point sets. Two vertices of the same layer are mapped to points with the same y-coordinate and an inner leaf and its corresponding split vertex are mapped to points with the same x-coordinate (see the highlighted orange circles). (For Theorem 8) The x-coordinate of $v = (6, 2)$ (green circle) can be bounded in terms of the difference between blue and red point in the axis-aligned rectangle with corners $(0, 0)$ and $\pi(v) = (M(v), \frac{v_x + v_y}{N}) = (\frac{10}{12}, \frac{8}{12})$. The rectangle contains 11 blue points and 3 red ones.

Proof. By definition of the auxiliary function, two leaves do not have the same mapping. Thus, it remains to show that the walk of two different split vertices cannot end at the same leaf. Imagine doing the walk backwards: start at any leaf, walk towards the origin and stop as soon as you reach a split vertex by traversing its non-preferred edge. Since each split vertex has exactly two children, it follows that exactly one leaf will stop at each split vertex. The exceptional case is the leaf $(N, 0)$, from which walking backwards to the origin is a horizontal path and the path does not contain any non-preferred edge. That is, in the inverse walk we follow preferred edges until we reach a non-preferred edge. This is equivalent to starting at a split vertex and follow the non-preferred edge once and continue with the preferred edges, which is the exact definition of our auxiliary function.   ◀

2.2   Transforming the tree into a pointset

With the auxiliary function M we can define the mapping between a weak CDR into a bicolored pointset in the unit square. For any vertex $v = (v_x, v_y) \in \mathcal{G}_N^+$ we define its *transformation* as $\pi(v) = (M(v), \frac{v_x + v_y}{N})$. Given any weak CDR, we look at the tree T it defines in $\mathcal{G}_N^+$. Each vertex $v \in \mathcal{D}$ creates a red point $\pi(v)$ and each split vertex $w \in \mathcal{S}$ creates a blue point $\pi(w)$ (note that we do not transform the boundary leaves in L_N into points). We define the *mapping* of T as the union of the sets $R = \{\pi(v) \colon v \in \mathcal{D}\}$ and $B = \{\pi(v) \colon v \in \mathcal{S}\}$ (see Figure 3, right). Note that the two sets depend on the tree T (and thus $R = R(T)$ and $B = B(T)$). From now on we assume that T is fixed, and thus we simplify the notation for ease of reading. For any set P of points in the unit square and $x, y \in [0, 1]$ let $P[x, y]$ be the number of points in $P \cap [0, x] \times [0, y]$.

▶ **Lemma 7.** *For any weak CDR T in $\mathcal{G}_N^+ \subset \mathbb{Z}^2$ and $n < N$, the red and blue points on the horizontal line $y = n/N$ alternate in color starting and ending with a blue point. In particular, we have $B[1, n/N] - R[1, n/N] = n + 1$.*

Proof. For the first statement we observe that only points that lie in L_n will have y-coordinates equal to n/N. Moreover, since L_{n+1} has one more vertex than L_n, each diagonal must have exactly one more split vertex than inner leaves. Indeed, Chun et al. showed that in proper CDRs each diagonal has exactly one split vertex (and of course, zero inner leaves).

Now we need to show that split vertices and inner leaves appear alternatingly on the diagonal line. Consider two consecutive split vertices $u, v \in L_n$ such that $u_x < v_x$. By definition of split, the edges $e_u = (u_x, u_y)(u_x + 1, u_y)$ and $e_v = (v_x, v_y)(v_x, v_y + 1)$ are all in T. Observe that there are $v_x - u_x - 1$ vertices in L_n and $v_x - u_x - 2$ vertices in L_{n+1} between e_u and e_v. Since two different vertices of L_n cannot connect to the same vertex of L_{n+1}, one of them will not reach L_{n+1}. That vertex will be an inner leaf and will be between u and v as claimed.

That is, the blue pointset has one more point than the red pointset in each horizontal line $y = i/N$. Summing up the differences from $i = 0$ to n, we get that in total there are $n + 1$ additional blue points $p = (x, y)$ with $y \le n/N$. ◀

With the above observations we can now state the main relationship between the weak CDR and its mapped pointset. For any vertex $v \in L_n$, its path to the origin splits the tree into two portions. Consider the portion of the tree up to L_n that is above the path from v to the origin. In L_0, the subtree contains a single vertex (the root) whereas at the diagonal L_n contains $v_x + 1$ vertices. Since the number of leaves grows with split vertices and shrinks with inner leaves, this means that in the portion of the tree that we are looking at, the difference between split vertices and inner leaves must be v_x, see Figure 3. Note that if the two children of a split vertex (e.g., $(5, 0)$ in Figure 3) are not in the same portion, the number of leaves does not grow with that split vertex. However, these split vertices may be still contained in the rectangle that we consider in the mapped pointset. This is the reason why we do not have an equality in Theorem 8.

▶ **Theorem 8.** *For any vertex $v \in \mathcal{G}_N^+$ it holds that $B[M(v), \frac{v_x + v_y}{N}] - R[M(v), \frac{v_x + v_y}{N}] - 2 \le v_x \le B[M(v), \frac{v_x + v_y}{N}] - R[M(v), \frac{v_x + v_y}{N}]$.*

Proof. We split the proof into two auxiliary lemmas.

▶ **Lemma 9.** *Let $v \in L_n$ be a split vertex such that $v_x < n$. If $M(v) < M(\gamma(v))$ the rectangle $[M(v), M(\gamma(v))] \times [0, \frac{n-1}{N}]$ contains exactly one point, which is blue and has $M(\gamma(v))$ as x-coordinate. If $M(\gamma(v)) < M(v)$ the rectangle $[M(\gamma(v)), M(v)] \times [0, \frac{n-1}{N}]$ contains exactly one point, which is blue and has $M(\gamma(v))$ as x-coordinate. When $v = (n, 0) \in L_n$ the rectangle $[M(v), M(\gamma(v))] \times [0, \frac{n-1}{N}]$ is empty.*

Proof. We first consider the case of $M(v) < M(\gamma(v))$. When we keep following from v to the preferred subtree, we end up in a leaf, called ℓ. By definition of M we have $M(\ell) = M(\gamma(v))$. Since $v_x < n$ we have $M(\gamma(v)) \ne 1$. By Lemma 6 there is a unique split vertex $s \in \mathcal{S}$ such that $M(s) = M(\ell)$. This split vertex is below layer L_n (indeed, we reach L_n from ℓ by following only preferred edges and the inverse walk has to stop when we traverse a non-preferred edge of s) and therefore s is transformed to a blue point in the rectangle. Now let s' be a split vertex which is mapped to a blue point in the rectangle. We will show that $s' = s$. Let ℓ' be the unique leaf such that $M(\ell') = M(s')$. Consider first the case in which ℓ' is below layer L_n (that is, $\ell'_x + \ell'_y < n$). Then let v' be the vertex on $dig(o, v)$ and $L_{\ell'_x + \ell'_y}$. If $\ell'_x < v'_x$ (resp. $v'_x < \ell'_x$) then Lemma 5 implies that $M(\ell') < min(T_{v'}) \le M(v)$ (resp. $M(\gamma(v)) \le max(T_{v'}) < M(\ell')$). This would be a contradiction to s' being mapped to a blue point in the rectangle.

It remains to consider the case in which ℓ' is above layer L_n. Define ℓ'' to be the vertex on $dig(o, \ell')$ and L_n. Lemma 5 implies that $\ell'' = v$ (otherwise we have either $M(\ell') < M(v)$ or $M(\gamma(v)) < M(\ell')$ which would again be a contradiction). Recall that there is only one split vertex whose walk to its corresponding leaf through preferred subtrees passes through v. Hence $s' = s$ and there is exactly one blue point in the rectangle.

We now show that there cannot be any red point either. Indeed, recall that for every red point there is a blue point with the same x-coordinate and smaller y-coordinate because for each inner leaf ℓ there is a unique split vertex s defined by the walk from s to ℓ such that $M(\ell) = M(s)$. From the previous argument, we know that s with $M(s) = M(\gamma(v))$ is mapped to the only one blue point in the rectangle and its corresponding leaf ℓ defined by the walk is above L_n. Hence, even if ℓ is an inner leaf, the mapped red point is not in the rectangle. Moreover, there cannot be any other red point in the rectangle (since it would imply that the corresponding blue point would also be in and we already ruled out this case).

In the same way we can also prove that if $M(\gamma(v)) < M(v)$ the rectangle $[M(\gamma(v)), M(v)] \times \left[0, \frac{n-1}{N}\right]$ contains exactly one point, which is blue and has $M(\gamma(v))$ as x-coordinate. If $v_x = n$ then ℓ as defined above is the leaf $(N, 0)$ and $M(\ell) = 1$. Lemma 6 implies that there is no split vertex s with $M(s) = 1$. ◀

▶ **Lemma 10.** *For any vertex $v \in \mathcal{G}_N^+$ it holds that*

$$v_x - B\left[M(v), \frac{v_x + v_y - 1}{N}\right] + R\left[M(v), \frac{v_x + v_y - 1}{N}\right] + 1 \in \{0, 1\}. \tag{1}$$

Proof. We first prove by induction over n that $\forall n \in \{0, 1, ..., N\}$ the following statement holds.

$$\{M(\gamma(p))|p \in L_n\}$$
$$= \left\{x \in [0, 1] : \left|B \cap \{x\} \times \left[0, \frac{n-1}{N}\right]\right| - \left|R \cap \{x\} \times \left[0, \frac{n-1}{N}\right]\right| = 1\right\} \cup \{1\}. \tag{2}$$

The quantity $|B \cap \{x\} \times [0, \frac{n-1}{N}]| - |R \cap \{x\} \times [0, \frac{n-1}{N}]|$ counts the difference between the number of blue points and red points on the vertical segment with x-coordinate x and length $\frac{n-1}{N}$. Because of Lemma 6 we know that each split vertex shares the same value with a leaf in the auxiliary function M. If the leaf is an inner leaf, both blue (split) and red (inner) points lie on the same unit segment $\{x\} \times [0, 1]$. Otherwise, there is only one blue point on $\{x\} \times [0, 1]$ because $M(p)$ for $p \in L_N$ are all different. Hence the quantity $|B \cap \{x\} \times [0, \frac{n-1}{N}]| - |R \cap \{x\} \times [0, \frac{n-1}{N}]|$ can either be 0 or 1.

The base case $n = 0$ trivially holds. We have $\{M(\gamma(p))|p \in L_0\} = \{1\}$ and

$$B \cap \{x\} \times \left[0, \frac{n-1}{N}\right] = R \cap \{x\} \times \left[0, \frac{n-1}{N}\right] = \emptyset$$

We assume that Section 2.2 holds for layer L_n and we prove that it also holds for L_{n+1}. We distinguish 3 cases for any vertex q in layer L_n.

- If q has degree 2 then q and its child $r \in L_{n+1}$ are mapped by $M \circ \gamma$ to the same value. Moreover q does not create any vertex in the set B nor R.
- If q is an inner leaf, then the value $M(\gamma(q))$ will not appear in $\{M(\gamma(p))|p \in L_{n+1}\}$ any more. The value $M(\gamma(q))$ also disappears in

$$\left\{x \in [0, 1] : \left|B \cap \{x\} \times \left[0, \frac{n}{N}\right]\right| - \left|R \cap \{x\} \times \left[0, \frac{n}{N}\right]\right| = 1\right\} \cup \{1\}.$$

because q created a red point in R with the coordinates $(M(\gamma(q)), \frac{n}{N}) = (M(q), \frac{n}{N})$.

- If q is a split vertex, then the value $M(\gamma(q))$ will stay in $\{M(\gamma(p))|p \in L_{n+1}\}$. Moreover $\{M(\gamma(p))|p \in L_{n+1}\}$ contains the additional value $M(q)$. The value $M(q)$ also appears in

$$\left\{ x \in [0,1] : \left| B \cap \{x\} \times \left[0, \frac{n}{N}\right] \right| - \left| R \cap \{x\} \times \left[0, \frac{n}{N}\right] \right| = 1 \right\} \cup \{1\}$$

because q creates a blue point in B with the coordinates $(M(q), \frac{n}{N})$.

Hence Section 2.2 holds.

Let v be a vertex in layer L_n, i.e. $n = v_x + v_y$. By Lemma 5 we know that a vertex $u \in L_n$ with $u_x < v_x$ satisfies $M(\gamma(u)) < M(\gamma(v))$. By Lemma 5 we also know that a vertex $w \in L_n$ with $v_x < w_x$ satisfies $M(\gamma(v)) < M(\gamma(w))$. Hence the number of vertices in layer L_n with smaller x-coordinate than that of v is exactly the number of vertices which are mapped by $M \circ \gamma$ to a smaller value than that of v. If $v_x < n$:

$$v_x = |\{u \in L_n|u_x < v_x\}| \overset{\text{Lemma } 5}{=} |\{u \in L_n|M(\gamma(u)) < M(\gamma(v))\}|$$
$$= |\{u \in L_n|M(\gamma(u)) \leq M(\gamma(v))\}| - 1$$
$$\overset{(2)}{=} B\left[M(\gamma(v)), \frac{n-1}{N}\right] - R\left[M(\gamma(v)), \frac{n-1}{N}\right] - 1$$
$$\overset{\text{Lemma } 9}{=} \begin{cases} B[M(v), \frac{n-1}{N}] - R\left[M(v), \frac{n-1}{N}\right] - 1 & \text{if } M(\gamma(v)) \leq M(v) \\ B[M(v), \frac{n-1}{N}] - R\left[M(v), \frac{n-1}{N}\right] & \text{if } M(v) < M(\gamma(v)) \end{cases}$$

If $v_x = n$ then:

$$v_x = |\{u \in L_n|M(\gamma(u)) \leq M(\gamma(v))\}| - 1 \overset{(2)}{=} B\left[M(\gamma(v)), \frac{n-1}{N}\right] - R\left[M(\gamma(v)), \frac{n-1}{N}\right]$$
$$\overset{\text{Lemma } 9}{=} B\left[M(v), \frac{n-1}{N}\right] - R\left[M(v), \frac{n-1}{N}\right] \qquad \blacktriangleleft$$

By Lemma 7, the red and blue points on the line $y = v_x + v_y$ alternate in color starting and ending with a blue point. Hence, any interval $[0, x]$ on the line $y = v_x + v_y$ contains at most one more blue points. Therefore, $B[M(v), \frac{v_x+v_y}{N}] - R[M(v), \frac{v_x+v_y}{N}] - (B[M(v), \frac{v_x+v_y-1}{N}] - R[M(v), \frac{v_x+v_y-1}{N}])$ is at most one. Lemmas 10 and 7 directly imply Theorem 8. $\blacktriangleleft$

3 Bichromatic discrepancy

Let R and B be a set of red and blue points in the unit square, respectively. Let $r = |R|$ and $b = |B|$, and further assume that $b > r$. Let $m = b - r$ (which is positive since $b > r$). For any set P of points in the unit square and $x, y \in [0, 1]$ let $P[x, y]$ be the number of points in $P \cap [0, x] \times [0, y]$.

For any two sets R and B and real numbers $x, y \leq 1$ we define the discrepancy of R and B at (x, y) as

$$D_{R,B}(x, y) = (b - r)xy - (B[x, y] - R[x, y]). \tag{3}$$

The *discrepancy* of R and B is simply defined as $D_{R,B}^* = \max_{(x,y) \in [0,1]^2} |D_{R,B}(x, y)|$ (i.e., the highest discrepancy we can achieve among all possible rectangles).

▶ **Theorem 2** (Two colors discrepancy). *For any set R and B of points such that $|B| > |R|$ it holds that $D_{R,B}^* = \Omega\left(\frac{(|B|-|R|) \cdot \log(|B|+|R|)}{|B|+|R|}\right)$.*

Note that if we set $R = \emptyset$ we get the classic two dimensional discrepancy result for which there are several proofs (see [12] for a detailed survey). In order to extend the bound for the case of $R \neq \emptyset$, we make minor changes to Schmidt's proof [13]. We start by using an auxiliary function G (defined below) and combining it with the trivial inequality

$$\int_{(x,y)\in[0,1]^2} D_{R,B}(x,y)G(x,y)dxdy \leq \max_{(x,y)\in[0,1]^2}|D_{R,B}(x,y)| \int_{(x,y)\in[0,1]^2}|G(x,y)|dxdy$$

to obtain $D_{R,B}^* = \max_{(x,y)\in[0,1]^2}|D_{R,B}(x,y)| \geq \frac{\int D_{R,B}G}{\int |G|}$.

Note that for simplicity in the notation we removed the integration limits. Our definition of G is identical to the one used by Schimdt: Let $m = \lceil \log_2(b+r) \rceil + 1$ and observe that, by definition of m we have $2(b+r) \leq 2^m \leq 4(b+r)$. For any $j \in \{0, \ldots, m\}$ we define function $f_j : [0,1]^2 \to \{-1,0,1\}$ as follows: subdivide the unit square with 2^j equally spaced vertical lines and 2^{m-j} horizontal lines.

For any value of j we subdivide the unit square into rectangles of area 2^{-m} (larger values of j will result in thinner but wider rectangles). Let A be a rectangle of subdivision associated to f_j. We define f_j within the rectangle to be 0 if A contains any point of $R \cup B$. If A does not have neither red nor blue points, we further subdivide it into four congruent quadrants. The function value of f_j is equal to 1 in the upper right and lower left quadrants, and -1 in upper left and lower right quadrants (see a visual representation of f_j in [12], page 173).

Then, we define G as $G = (1 + cf_0)(1 + cf_1)\ldots(1 + cf_m) - 1$, where $c > 0$ is a small constant (whose value will be chosen afterwards). Note that G can also be expressed as $G = G_1 + \ldots G_m$, where $G_k = c^k \sum_{0 \leq j_1 \leq \ldots \leq j_m \leq m} f_{j_1} f_{j_2} \cdots f_{j_k}$.

Schmidt showed that $\int |G| \leq 2$ (regardless of the value of m). Thus, we now focus in giving an upper bound for $\int D_{R,B}G$.

▶ **Lemma 11.** *There exists a constant c_1 such that $\int D_{R,B}G_1 \geq cc_1 \frac{b-r}{b+r}\log(b+r)$.*

Proof. By definition of G_1 we have $\int D_{R,B}G_1 = c\sum_{j=0}^{m} \int D_{R,B}f_j$. Thus, it suffices to show that for any value of j it holds that $\int D_{R,B}f_j \geq c'\frac{b-r}{b+r}$ (for some other constant $c' > 0$).

Recall that, when defining f_j, we subdivided the unit square into at least $2(b+r)$ rectangles. For the rectangles that contain at least one point of $R \cup B$, f_j is set to zero, and thus they do not contribute to the integral. Since we have $b+r$ many points, we know that there must exist at least $b+r$ rectangles that do not contain any point of R or B. Let A be any such rectangle, and let $A_{SW}, A_{NW}, A_{SE}, A_{NE}$ be the four subquadrants of A (where the subindex refers to the cardinal position of the quadrant). Recall that f_j is equal to 1 for any point of $A_{SW} \cup A_{NE}$ and -1 for points of $A_{SE} \cup A_{NW}$.

Let w and h be vectors defined by the horizontal and vertical sides of A_{SW}, respectively. Observe that their lengths are 2^{-j-1} and 2^{j-m-1}, respectively. Then, we have

$$\int_A f_j D_{R,B}$$
$$= \int_{A_{SW}} D_{R,B} - \int_{A_{NW}} D_{R,B} + \int_{A_{NE}} D_{R,B} - \int_{A_{SE}} D_{R,B}$$
$$= \int_{A_{SW}} [D_{R,B}(x,y) + D_{R,B}(x+\mathsf{w},y+\mathsf{h}) - D_{R,B}(x,y+\mathsf{h}) - D_{R,B}(x+\mathsf{w},y)]dxdy.$$

If we apply the definition of $D_{R,B}$ (Eq. (3)) to the four terms inside the integral we get

$$
\begin{aligned}
\int_A f_j D_{R,B} \;=\; & \int_{A_{SW}} ((b-r)[xy + (x+\mathsf{w})(y+\mathsf{h}) - x(y+\mathsf{h}) - (x+\mathsf{w})y]) dxdy \\
& - \int_{A_{SW}} (B[x,y] + B[x+\mathsf{w}, y+\mathsf{h}] - B[x,y+\mathsf{h}] - B[x+\mathsf{w}, y]) dxdy \\
& + \int_{A_{SW}} (R[x,y] + R[x+\mathsf{w}, y+\mathsf{h}] - R[x,y+\mathsf{h}] - R[x+\mathsf{w}, y]) dxdy.
\end{aligned}
$$

Observe that we are integrating twice positively and twice negatively over almost identical functions. In fact, the terms of the first integral all cancel out except along the rectangle $[x, x+\mathsf{w}) \times [y, y+\mathsf{h})$. Similarly, when we look at the second and third terms, the contribution of any point in $R \cup B$ is cancelled out unless it is in the rectangle $[x, x+\mathsf{w}) \times [y, y+\mathsf{h})$. However, by definition of A there are no such points. Thus, we obtain

$$
\int_A f_j D_{R,B} = \int_{A_{SW}} (b-r)\mathsf{w} \cdot \mathsf{h} \, dxdy = \int_{A_{SW}} (b-r)2^{-m-2} dxdy = (b-r)2^{-2m-4}
$$

That is, when we integrate $f_j D_{R,B}$ over a rectangle A containing no point of $R \cup B$, the result is $(b-r)2^{-2m-4}$. We know that there are at least $b+r$ rectangles not containing points of $R \cup B$, thus their contribution is at least $\frac{(b+r)(b-r)}{2^{2m+4}} = \frac{(b+r)}{2^m} \frac{(b-r)}{16 \cdot 2^m} \geq \frac{1}{4} \frac{(b-r)}{16 \cdot 4(b+r)} = \Omega(\frac{b-r}{b+r})$. ◄

▶ **Lemma 12.** *There exists a constant c_2 such that $\sum_{k=2}^m \int D_{R,B} G_k \leq c^2 c_2 \frac{b-r}{b+r} \log(b+r)$.*

Proof. Recall that $G_k = c^k \sum_{0 \leq j_1 < j_2 < \ldots < j_k \leq m} f_{j_1} \ldots f_{j_k}$. Fix any valid set of indices and consider the value of $\int f_{j_1} \ldots f_{j_k} D_{R,B}$.

As shown in [12], function $f_{j_1} \ldots f_{j_k}$ is largely defined by f_{j_1} and f_{j_k}. Indeed, if we overlay the rectangular partition defined by functions $f_{j_1}, \ldots, f_{j_k}$ we obtain a grid of rectangles whose width is 2^{-j_k} and height $2^{-(m-j_1)}$. In each of these rectangles, the function is zero (if any of the rectangles associated to the f_{j_i} functions contains a point of $R \cup B$), or is further subdivided into four equal sized quadrants and in each one it is $+1$ or -1 alternatively.

Let A be one of the rectangles of the refined grid. As shown in Lemma 11, we have that

$$
\int_A f_{j_1} \ldots f_{j_k} D_{R,B} = \tau (b-r)2^{-2(m+j_k-j_1)-4},
$$

where $\tau \in \{-1, 1\}$. This extra term appears because the product of the different functions involved can change the sign of each of the four quadrants. In any case, we have $\int_A f_{j_1} \ldots f_{j_k} D_{R,B} \leq (b-r)2^{-2(m+g)-4}$ where $g = j_k - j_1$.

By the way the grid is constructed, there are $2^{m-j_1} \times 2^{j_k} = 2^{m+g}$ many rectangles, and thus we conclude that $\int f_{j_1} \ldots f_{j_k} D_{R,B} \leq (b-r)2^{-m-g-4}$. In order to obtain a bound $\int D_{R,B} G_k$ we sum over all possible indices.

$$
\int D_{R,B} G_k = c^k \sum_{0 \leq j_1 < \ldots < j_k \leq m} \int f_{j_1} \ldots f_{j_k} D_{R,B} \leq \frac{c^k(b-r)}{2^{m+4}} \sum_{0 \leq j_1 < \ldots < j_k \leq m} 2^{-(j_k-j_1)}.
$$

Note that in the sum, the indices $j_2, \ldots j_{k-1}$ do not matter. Thus, we group the terms by the gap between the indices j_1 and j_k (say, if $j_1 = 3$ and $j_k = 7$ the gap is 4). Note that the minimum gap is at least $k-1$ (since otherwise we do not have enough space to choose the $k-2$ indices in between) and at most m. Once we have a gap of g there are $m-g$ options for index j_1.

$$\int D_{R,B} G_k \leq \frac{c^k(b-r)}{2^{m+4}} \sum_{g=k-1}^{m} \sum_{j_1=0}^{m-g} \sum_{j_1<j_2<\ldots<j_{k-1}<j_1+g} 2^{-g}$$

$$= \frac{c^k(b-r)}{2^{m+4}} \sum_{g=k-1}^{m} \sum_{j_1=0}^{m-g} \binom{g-1}{k-2} 2^{-g} \leq \frac{c^k(b-r)m}{2^{m+4}} \sum_{g=k-1}^{m} \binom{g-1}{k-2} 2^{-g}.$$

In order to upper bound the sum over all G_k, we first reorder the summation order.

$$\sum_{k=2}^{m} \int D_{R,B} G_k \leq \sum_{k=2}^{m} \frac{c^k(b-r)m}{2^{m+4}} \sum_{g=k-1}^{m} \binom{g-1}{k-2} 2^{-g}$$

$$= \frac{(b-r)m}{2^{m+4}} \sum_{g=1}^{m} 2^{-g} c^2 \sum_{k=2}^{g+1} \binom{g-1}{k-2} c^{k-2}$$

$$= \frac{(b-r)m}{2^{m+4}} \sum_{g=1}^{m} 2^{-g} c^2 (1+c)^{g-1}$$

$$= \frac{(b-r)mc^2}{2^{m+5}} \sum_{g=1}^{m} \left(\frac{1+c}{2}\right)^{g-1}.$$

The sum contains the first terms of the geometric sum $\sum_{g=1}^{\infty} \left(\frac{1+c}{2}\right)^{g-1} \leq \frac{2}{1-c}$ (for any $c < 1$). In particular, if we set $c \leq 1/2$ we can upper bound the partial sum by 4. Recall that $m = \Theta(\log(b+r))$ and $2^m = \Theta(b+r)$. Thus, the lemma is proven. ◀

▶ **Corollary 13.** *There exists a constant $\kappa > 0$ such that $\int D_{R,B} G \geq \kappa \left(\frac{(b-r)\cdot\log(b+r)}{b+r}\right)$.*

Proof. Apply the inequality $\int(A+B) \geq \int A - \int |B|$ and Lemmas 11 and 12 to obtain:

$$\int D_{R,B} G = \int D_{R,B} G_1 + \sum_{k=2}^{m} \int D_{R,B} G_k \geq c(c_1 - cc_2) \left(\frac{(b-r)\cdot\log(b+r)}{b+r}\right)$$

Note that Lemmas 11 and 12 holds for any value of c such that $c \in (0, 1/2]$. By choosing a sufficiently small value of c (say, $c = \min\{\frac{1}{2}, \frac{c_1}{2c_2}\}$) we obtain $\int D_{R,B} G \geq \frac{cc_1}{2} \left(\frac{(b-r)\cdot\log(b+r)}{b+r}\right)$. ◀

This completes the proof of Theorem 2.

When $R = \emptyset$, it would be expected that we need to distribute the blue points uniformly in the unit square to have a low discrepancy. Indeed, it is also held for the red points. The following theorem implies that even if there are many red points, but the red points are concentrated in the lower half of the unit square, the discrepancy cannot be reduced. For simplicity, we only show a special case of how the discrepancy is depended on the points in $[0, 1] \times [1/2, 1]$, which is good enough for our purpose in Section 5. Notice that the same argument can be applied in a more general case.

▶ **Theorem 14.** *For any set R and B of points in the unit square such that $|R| = r$, $|B| = b$ and $b > r$. Let r_2 and b_2 be the number of red and blue points in $[0, 1] \times [1/2, 1]$ respectively. It holds that*

$$D_{R,B}^* = \Omega\left(\frac{(b_2 - r_2)\cdot\log(b_2 + r_2)}{b_2 + r_2}\right).$$

Proof. Let R_2 and B_2 be the set of red and blue points in $[0,1] \times [1/2, 1]$ respectively. Consider the upper half of the unit square $[0,1] \times [1/2, 1]$ and rescale the vertical length to be 1. By Theorem 2, there exists a point $(x, 2y)$ such that $|D_{R_2, B_2}(x, 2y)| = |2xy(b_2 - r_2) - (B_2[x, 2y] - R_2[x, 2y])| \geq 2c \left(\frac{(b_2 - r_2) \cdot \log(b_2 + r_2)}{b_2 + r_2} \right)$ for some constant c.

Then, we map the point $(x, 2y)$ back to a point $(x, 1/2 + y)$ in the original unit square. We will show that either $D_{R,B}(x, 1/2 + y)$ or $D_{R,B}(x, 1/2 - \epsilon)$ would give us the desired lower bound, where ϵ is an arbitrarily small constant such that rectangle $[0,1] \times [0, 1/2 - \epsilon]$ only contains $B \setminus B_2$ and $R \setminus R_2$.

If $|D_{R,B}(x, 1/2 + y)| \geq c \left(\frac{(b_2 - r_2) \cdot \log(b_2 + r_2)}{b_2 + r_2} \right)$, we are done.

If $|(b - r)/2 - (b_2 - r_2)| \geq c/4 \left(\frac{(b_2 - r_2) \cdot \log(b_2 + r_2)}{b_2 + r_2} \right)$, the proof is also done. Because

$$
\begin{aligned}
& |D_{R,B}(1, 1/2 - \epsilon)| \\
\overset{(3)}{=} \ & |(b - r)(1/2 - \epsilon) - (B[1, 1/2 - \epsilon] - R[1, 1/2 - \epsilon])| \\
= \ & |(b - r)(1/2 - \epsilon) - (b - r - (b_2 - r_2))| \\
= \ & |(b - r)/2 - (b_2 - r_2) - (b - r)\epsilon| \\
> \ & c/8 \left(\frac{(b_2 - r_2) \cdot \log(b_2 + r_2)}{b_2 + r_2} \right).
\end{aligned}
$$

Suppose that the two cases do not hold, we have $|D_{R,B}(x, 1/2 + y)| < c \left(\frac{(b_2 - r_2) \cdot \log(b_2 + r_2)}{b_2 + r_2} \right)$ and $|(b - r)/2 - (b_2 - r_2)| < c/4 \left(\frac{(b_2 - r_2) \cdot \log(b_2 + r_2)}{b_2 + r_2} \right)$. Let $R_1 = R \setminus R_2$ and $B_1 = B \setminus B_2$, which are inside the rectangle $[0,1] \times [0, 1/2 - \epsilon]$. Consider

$$
\begin{aligned}
& D_{R,B}(x, 1/2 + y) \\
= \ & (b - r)x(1/2 + y) - (B[x, 1/2 + y] - R[x, 1/2 + y]) \\
= \ & (b - r)x(1/2 + y) - (B_2[x, 1/2 + y] - R_2[x, 1/2 + y] + B_1[x, 1/2 - \epsilon] - R_1[x, 1/2 - \epsilon]) \\
= \ & (b - r)x(1/2 - \epsilon) - (B_1[x, 1/2 - \epsilon] - R_1[x, 1/2 - \epsilon]) + (b - r)x\epsilon \\
& + (b - r)xy - (B_2[x, 1/2 + y] - R_2[x, 1/2 + y]) \\
= \ & D_{R,B}(x, 1/2 - \epsilon) + (b - r)xy - (B_2[x, 1/2 + y] - R_2[x, 1/2 + y]) + (b - r)x\epsilon \\
> \ & D_{R,B}(x, 1/2 - \epsilon) + 2(b_2 - r_2)xy - (B_2[x, 1/2 + y] - R_2[x, 1/2 + y]) \\
& - c/2 \left(\frac{(b_2 - r_2) \cdot \log(b_2 + r_2)}{b_2 + r_2} \right) + (b - r)x\epsilon \\
= \ & D_{R,B}(x, 1/2 - \epsilon) + D_{R_2, B_2}(x, 2y) - c/2 \left(\frac{(b_2 - r_2) \cdot \log(b_2 + r_2)}{b_2 + r_2} \right) + (b - r)x\epsilon
\end{aligned}
$$

The first inequality is given by $b - r > 2(b_2 - r_2) - c/2 \left(\frac{(b_2 - r_2) \cdot \log(b_2 + r_2)}{b_2 + r_2} \right)$. Since $|D_{R,B}(x, 1/2 + y)| < c \left(\frac{(b_2 - r_2) \cdot \log(b_2 + r_2)}{b_2 + r_2} \right)$ and $|D_{R_2, B_2}(x, 2y)| \geq 2c \left(\frac{(b_2 - r_2) \cdot \log(b_2 + r_2)}{b_2 + r_2} \right)$, we can conclude that $|D_{R,B}(x, 1/2 - \epsilon)| = \Omega(\frac{(b_2 - r_2) \cdot \log(b_2 + r_2)}{b_2 + r_2})$. ◀

4   Lower bound for two dimensional weak CDRs

Before giving the proof of Theorem 3, we recall that a proof for a proper CDR (i.e., one without inner leaves) was given in [8]. Our proof follows the same spirit, so we first give an overview of their proof and describe what changes when we introduce inner leaves.
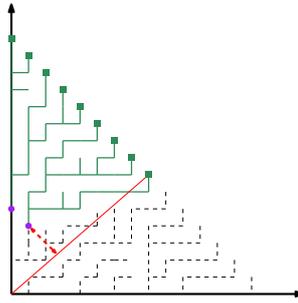
■ **Figure 4** Illustration of why the two sets I (purple disks) and $\mathcal{L}(I)$ (green squares) should have proportional sizes. If the size of $\mathcal{L}(I)$ grows drastically (as shown in the figure), the point of the highest x-coordinate in $\mathcal{L}(I)$ must make a significant detour to pass through I, causing a large error. A similar effect happens if the size of $\mathcal{L}(I)$ is comparatively small.

▶ **Lemma 15.** *Given a CDR, a point $p = (x, y) \in L_N$, and an integer $n < N$, let $p' = (x', y') \in L_n$ be the unique point of L_n that is in $dig(o, p)$. The Hausdorff error of the CDR is at least $|x' - x \cdot \frac{n}{N}|$.*

Proof. This result was shown by Chun et al. [8] (Lemma 3.5, in Cases 1 and 2). We give the proof for completeness. Consider the L-infinity ball of radius $|x' - x \cdot \frac{n}{N}|$ centered at $p \cdot \frac{n}{N}$. By construction, this ball contains p' in its boundary. Because of the monotonicity axiom, no vertex of $dig(o, p)$ can be in the interior of the ball. In particular, when measuring the Hausdorff distance of point $p \cdot \frac{n}{N} \in \overline{op}$ we get an error of at least $|x' - x \cdot \frac{n}{N}|$. ◀

Consider any point $p \in L_N$ and virtually sweep a line of slope -1 from the origin all the way to L_N. During the sweep, the intersection between the diagonal line and either the Euclidean segment $\overline{op}$ or the digital one $dig(o, p)$ will be a point. Lemma 15 says that if we can find an instant of time for which two intersection points are at distance ∂ from each other, then the Hausdorff error of the whole CDR must be $\Omega(\partial)$ (see Figure 4).

In order to find this instant of time we see how much the subtrees grow. Consider a consecutive set of I vertices in some intermediate layer L_n. Let $\mathcal{L}(I)$ be the vertices of L_N whose digital path to the origin passes through some vertex of I. If the CDR has small error, we need $\mathcal{L}(I)$ to have roughly $\frac{N}{n}|I|$ many points. The difference between the expected number of vertices and $|\mathcal{L}(I)|$ combined with Lemma 15 will give a lower bound on the Hausdorff error.

Our proof follows the same spirit (transform the tree into a pointset, use discrepancy to find a subset with too many/too few children and use Lemma 15 to find a large error). Although all three steps follow the same spirit, they need major changes to account for the possibility of inner leaves.

The biggest change is how we map the tree. In proper CDRs each line has a unique split vertex and always extends to L_N. Thus, a region with a large number of split vertices directly implies a large error. In our setting, we could potentially have a region with many split vertices followed by a large number of inner leaves to cancel out the growth. This is why we need two major changes: first we now color the points red and blue depending on whether they are split vertices or inner leaves. We also introduce a second dimension to track when the children of a split vertex stop extending. Intuitively speaking, the x-coordinate of our mapping will be similar to the mapping done by Chun et al. [8] whereas the y-coordinate represents time. Thus, the difference in y-coordinates between red and blue points can be used to determine for how long are the two children of a split vertex alive (the longer the difference in y-coordinates, the further away that the two children extend).

We now use the mapping of Section 2 together with the two colors discrepancy (Theorems 2 and 14) to show a lower bound on the error of weak CDRs. The discrepancy result in Theorem 2 considers the points in the whole unit square. Due to some technical reasons, in Section 5 we will need a discrepancy result for the points in the *upper half* of the unit square instead (Theorem 14). The difference between the two theorems is just a constant factor and thus would have little implication. Here we use Theorem 14 and prove the result in terms of the number of inner leaves in the upper half. Specifically, we show the following result.

▶ **Theorem 3.** *For any $N \in \mathbb{N}$, any weak CDR defined on $\mathcal{G}_N^+ \subset \mathbb{Z}^2$ with κ_2 inner leaves between lines $x + y = \lceil N/2 \rceil$ and $x + y = N$ has $\Omega(\frac{N \log N}{N + \kappa_2})$ error.*

Proof. Given a weak CDR and its associated tree T, consider its transformation into the sets R and B of red and blue points defined by π. Let b_2 and r_2 be the numbers of blue and red points in the rectangle $[0,1] \times [1/2,1]$ respectively. By Lemma 7, we have $b_2 - r_2 = \lfloor N/2 \rfloor$. We apply the discrepancy result (Theorem 14) with $b_2 - r_2 = \lfloor N/2 \rfloor$ and $r_2 = \kappa_2$, and obtain that there exists $\alpha, \beta \in [0,1]$ such that $|B[\alpha, \beta] - R[\alpha, \beta] - N \cdot \alpha \cdot \beta| > c' \cdot \frac{N \cdot \log N}{N + \kappa_2}$.

We want to use Theorem 8 on the vertex of T whose image is (α, β). Naturally, such a vertex need not exist, but we will find one nearby whose associated discrepancy is also high. Let $n = \lfloor N \cdot \beta \rfloor$ and observe that $B[\alpha, \beta] = B[\alpha, \frac{n}{N}]$; indeed, by the way we transform points, their y-coordinates are of the form i/N. However, by definition of n we know that β is between n/N and $(n+1)/N$ and thus no point can lie in the horizontal strip $y \in (n/N, \beta]$ (by the same argument we also have $R[\alpha, \beta] = R[\alpha, \frac{n}{N}]$).

If we substitute β in the previous equation we get

$$\left| B\left[\alpha, \frac{n}{N}\right] - R\left[\alpha, \frac{n}{N}\right] - \alpha n \right| > c' \cdot \frac{N \log N}{N + \kappa_2} - 1 \geq c'' \cdot \frac{N \cdot \log N}{N + \kappa_2}$$

for a large enough N, $\kappa_2 \in O(N \log N)$ and for some $c'' > 0$. We get the additional 1 term because of the rounding in the definition of n.

Now we need to do a similar operation for α. Let $q_i = (i, n - i)$ be a vertex of L_n. By Lemma 5 the image of the auxiliary function $M(q_i)$ monotonically increases as i grows. Let $Q = \{q_i : M(q_i) \leq \alpha\}$ and $\alpha' = max_{q_i \in Q} M(q_i)$. Note that, by definition of the set Q, it trivially holds that $\alpha' \leq \alpha$.

▶ **Lemma 16.** $B[\alpha, \frac{n}{N}] - R[\alpha, \frac{n}{N}] = B[\alpha', \frac{n}{N}] - R[\alpha', \frac{n}{N}]$

Proof. The difference between the two rectangles is the rectangle Δ whose opposite corners are $(\alpha', 0)$ and $(\alpha, n/N)$, and one of the boundary $\overline{(\alpha', 0)(\alpha', \frac{n}{N})}$ is open. We claim that red and blue points are paired (sharing the same x-coordinate) in Δ (and thus, for each red point that we remove we are also removing a blue one). By Lemma 6, we know that all the blue points have different x-coordinates, so do red points. Hence, if there are red and blue points on the same vertical line, they must be the only pair in that vertical line. First notice that if there is a red point in Δ, there also exists a blue point in Δ with the same x-coordinate and below the red point. By the virtual walk that we define the auxiliary function, every split vertex is closer to the origin than the corresponding leaf. Hence, after the transformation π, if there is a red point, then there must exist a blue point with the same x-coordinate (by Lemma 6) and smaller y-coordinate. Then, we will show that if there is a blue point in Δ, there also exists a red point in Δ with the same x-coordinate.

Assume, for the sake of contradiction that there exists a blue point p in Δ such that there does not exist a red point q with the same x-coordinate as p in Δ. Let s be the split vertex whose image is p. By definition of the transformation π, the x-coordinate of p is $M(s)$, which
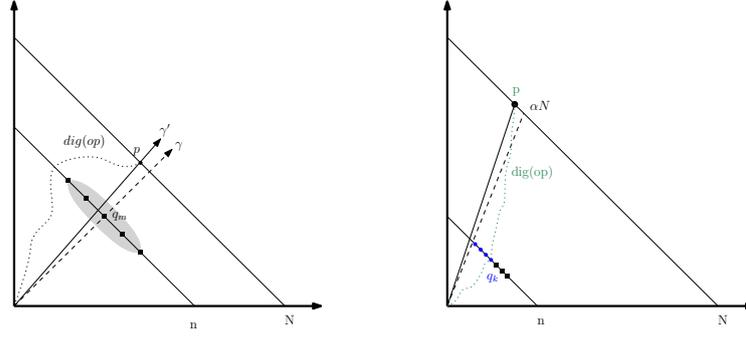
Figure 5 (left) When k is small we have $\Omega(\frac{N \log N}{N + \kappa_2})$ consecutive vertices in L_n that are not productive (shown as squares). In particular, the ray γ through the middle point must make a large detour. (right) When k is large, there is a digital path through q_k with a big detour.

is between α' and α. We apply Lemma 6 to find the unique leaf ℓ such that $M(s) = M(\ell)$. Since $\pi(\ell) \notin \Delta$, we have that $\ell_x + \ell_y > n$. Let m be the unique vertex of L_n that is in the path from s to ℓ. It follows that $\pi(m) = (M(\ell), \frac{n}{N}) \in \Delta$. This gives a contradiction with the definition of α', and thus implies that if there exists a blue point in Δ, then there also exists a red point in Δ with the same x-coordinate. ◀

Thus, given a pair (α, β) whose associated rectangle has high discrepancy, we have *snapped* it to the pair $(\alpha', \frac{n}{N})$ that defines another rectangle with high discrepancy. More importantly, by definition of Q, we know that $\pi(q_{|Q|-1}) = (\alpha', \frac{n}{N})$. Note that $q_{|Q|-1}$ need not be a split vertex or an inner leaf (and thus, $(\alpha', \frac{n}{N})$ may not be a point of $R \cup B$).

Let $b' = B[\alpha', \frac{n}{N}]$ and $r' = R[\alpha', \frac{n}{N}]$. If we apply Theorem 8 to point $q_{|Q|-1}$ we get that $b' - r' - 2 \le |Q| - 1 \le b' - r'$. This set Q is the one that makes the role of I in the proof overview: we know that vertices of Q are the ones that extend to cover all the vertices of L_N whose image is α' or less. As such, we would expect $|Q|$ to contain roughly $n\alpha'$ elements. However, the discrepancy result tells us that the size of Q is $c'' \frac{N \log N}{N + \kappa_2}$ units away from that value. We say that p is *productive* if some point of $T(p)$ is in L_N (this is equivalent to the fact that p can be extended to reach the boundary). Let $k \le b' - r' - 2$ be the biggest integer such that q_k is productive. Note that k is well defined because q_0 is always productive ($(0, n)$ always extends to $(0, N)$). The proof now considers a few cases depending on whether k is small or large (specifically, we say that k is *small* if $|Q| - 1 - k \ge (b' - r' - 2) - k > \frac{c''}{2} \cdot \frac{N \log N}{N + \kappa_2}$, *large* otherwise) and if Q contains too few or too many points.

k is small. Recall that we looked for the largest possible k (such that q_k is productive). Thus, if k is small, we have many points in layer L_n that are consecutive and not productive. In particular, none of the vertices in $q_{b'-r'-\lfloor \frac{c''}{2} \cdot \frac{N \log N}{N + \kappa_2} \rfloor}, \ldots, q_{b'-r'-2}$ are productive. Let $q_m = q_{b'-r'-\lfloor \frac{c''}{4} \cdot \frac{N \log N}{N + \kappa_2} \rfloor}$ (note that this point is surrounded by non-productive points in both sides along L_n).

Shoot a ray γ from o towards q_m. Let p be the vertex on L_N that is closest to γ. Observe that the $|| \cdot ||_\infty$ distance between γ and p is at most $1/2$. Let γ' be the ray shooting from o towards p. Similarly, the $|| \cdot ||_\infty$ distance between γ' and q_m is at most $1/2$ (see Figure 5, left).

We now apply Lemma 15 to $dig(o, p)$. We know that the Euclidean segment $\overline{op}$ is close to q_m. The digital segment must cross L_n and is far from q_m (the closest it can pass is either $q_{b'-r'-\lfloor \frac{c''}{2} \cdot \frac{N \log N}{N + \kappa_2} \rfloor - 1}$ or $q_{b'-r'-1}$). That is, we know that the intersection of $\overline{op}$

with the line $x + y = n$ is at most half a unit away from q_m. Similarly, the intersection with $dig(o, p)$ is at least $\lfloor \frac{c''}{4} \cdot \frac{N \log N}{N + \kappa_2} \rfloor$ from q_m. Thus, by triangle inequality the $|| \cdot ||_\infty$ distance between $dig(o, p)$ and $\overline{op}$ is at least $\lfloor \frac{c''}{4} \cdot \frac{N \log N}{N + \kappa_2} \rfloor - 3/2 \in \Omega(\frac{N \log N}{N + \kappa_2})$.

k is large and $b' - r' \geq n\alpha + c'' \cdot \frac{N \log N}{N + \kappa_2}$. Look at the x-coordinate of q_k. We know that Q has at least $b' - r' - 1 \geq n\alpha + c'' \cdot \frac{N \log N}{N + \kappa_2} - 1$ many elements, and k is among the productive vertices with the largest x-coordinate. In particular, the x-coordinate of q_k is at least $b' - r' - 2 \geq n\alpha + \frac{c''}{2} \cdot \frac{N \log N}{N + \kappa_2} - 2$.

Let p be the unique leaf of L_N such that $M(p) = M(q_k)$. We now apply Lemma 15 to $dig(o, p)$ at the line $x + y = n$. By definition of p, we have that $dig(o, p)$ passes through q_k. Now, by definition of Q, we know that $M(q_k) \leq \alpha$ and in particular the x-coordinate of p is at most αN (see Figure 5, right). Thus, the Euclidean segment $\overline{op}$ must intersect at a point whose x-coordinate is at most αn.

That is, when we look at the Euclidean and the digital segments along line $x + y = n$, the Euclidean crossing happens at x-coordinate at most αn. However, the x-coordinate of the digital crossing is at least $\alpha n + \frac{c''}{2} \cdot \frac{N \log N}{N + \kappa_2} - 1$. By Lemma 15 we conclude that the error must be $\Omega(\frac{N \log N}{N + \kappa_2})$ as claimed.

$b' - r' < n\alpha - c'' \cdot \frac{N \log N}{N + \kappa_2}$. This proof is very similar to the previous case. Consider the vertex $p = (\lfloor \alpha N \rfloor, N - \lfloor \alpha N \rfloor) \in L_N$ and apply Lemma 15 to $dig(o, p)$ and $\overline{op}$.

At line $x + y = n$ the Euclidean segment $\overline{op}$ passes through a point whose x-coordinate is $\lfloor \alpha N \rfloor \cdot \frac{n}{N} \geq \lfloor \alpha n \rfloor - 1$. By definition, $M(p) \leq \alpha$ and thus $dig(o, p)$ must pass through some vertex q of Q. In particular, the x-coordinate of q is at most $b' - r' < n\alpha - c'' \cdot \frac{N \log N}{N + \kappa_2}$, giving the $\Omega(\frac{N \log N}{N + \kappa_2})$ error and completing the proof of Theorem 3.   ◄

Note that if we use Theorem 2 instead, the same argument follows and we would get the following result.

▶ **Theorem 17.** *For any $N \in \mathbb{N}$, any weak CDR defined on $\mathcal{G}_N^+ \subset \mathbb{Z}^2$ with κ_1 inner leaves has $\Omega(\frac{N \log N}{N + \kappa_1})$ error.*

5   Lower bound for CDRs in high dimensions

We now use the lower bound of weak CDRs to obtain a lower bound for CDRs in three or higher dimensions. Consider the restriction of any d-dimensional CDR T to the $x_1 x_2$-plane (we call this restriction the $x_1 x_2$-*restriction of T* and denote it by $T_{x_1 x_2}$). Recall that the key observation is that $T_{x_1 x_2}$ is a (possibly weak) CDR and that any inner leaf in $T_{x_1 x_2}$ must extend in some x_i-direction in T for some $i \in [3..d]$. We have seen that $T_{x_1 x_2}$ needs to have a large number of inner leaves to have $o(\log N)$ error. In the following, we will show that a large number of inner leaves will cause constraints for $\mathbb{Z}^d$ and have an impact in the overall error of T.

We do a slight abuse of notation and use the same terms as in two dimensions. For simplicity of the notation, we assume that N is a positive even number. For any $n \leq N$, let $L_n = \{(x_1, x_2, \ldots, x_d) \in \mathcal{G}_N^+ : \sum_{i=1}^d x_i = n\}$. Given any CDR in $\mathcal{G}_N^+$, we consider the CDR as a tree rooted at the origin. Let $T(v)$ be the subtree rooted at v.

From Theorem 3, we already know that in order for $T_{x_1 x_2}$ to have sublogarithmic error we must have $\kappa_2 \in \omega(N)$ inner leaves. However, each inner leaf ties to a boundary leaf in L_N in d dimensions. In other words, the subtrees rooted at the vertices in $L_{N/2-1} \cap T_{x_1 x_2}$ must cover all these boundary vertices. We now observe that a weak CDR with inner leaves in the $x_1 x_2$-plane induces subtrees which are too big for the high dimensional proper CDR (See Figure 6).
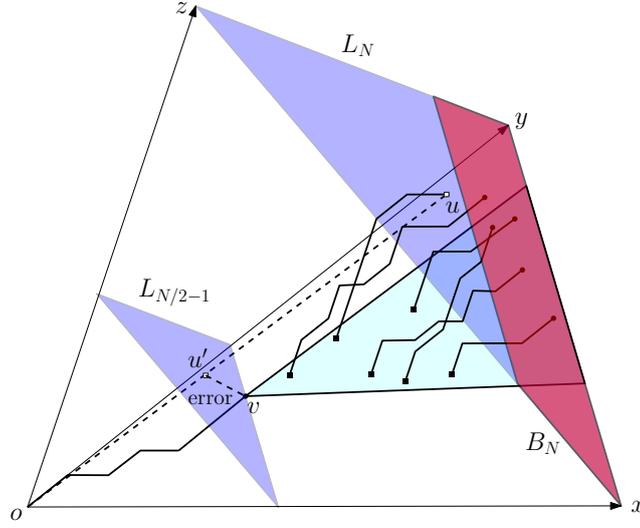
■ **Figure 6** Illustration of Lemmas 18 and 19: the red region represents the region of B_N. If we have lots of inner leaves in T_{xy}, it will have many descendants in the three dimensional CDR at layer L_N so that the height of the red region attempting to contain them is large. In particular, we can find a vertex v on the xy-plane such that v is on the $dig(o, u)$ and u is far away from the xy-plane. For simplicity, we show the Euclidean error between v and u', but we note that the proof argues under the $|| \cdot ||_\infty$ metric.

▶ **Lemma 18.** *Given any CDR in $\mathcal{G}_N^+$, let κ_2 be the number of inner leaves in $T_{x_1 x_2}$ between $L_{N/2}$ and L_N. There exists a vertex $v \in L_{N/2-1}$ such that $v_i = 0$ for $i = 3, \ldots, d$ and some boundary leaf $u \in T(v) \cap L_N$ has $u_j \geq (\kappa_2/N)^{\frac{1}{d-2}} - 1$ for some $j \in [3..d]$.*

Proof. The proof follows from a packing argument. Consider the set $V = \{(0, N/2 - 1, 0, \ldots, 0), (1, N/2 - 2, 0, \ldots, 0), \ldots, (N/2 - 1, 0, 0, \ldots, 0)\}$. Note that these vertices lie in the $x_1 x_2$-plane and thus are in $T_{x_1 x_2}$. Because they are the two dimensional equivalent of $L_{N/2-1}$, the union of their subtrees covers $T_{x_1 x_2}$ between $N/2$ and N. In this region we know that we have κ_2 many inner leaves, which will extend to L_N with the first step in the x_i-direction for some $i \in [3..d]$. Let Y_N be the extended vertices on L_N from these κ_2 inner leaves, i.e., $|Y_N| \geq \kappa_2$.

Let $B_N = \{(x_1, x_2, \ldots, x_d) \in \mathcal{G}_N^+ \colon \sum_{i=1}^d x_i = N, x_1 + x_2 < N$ and $\forall i \in [3..d], x_i < (\kappa_2/N)^{\frac{1}{d-2}} - 1\}$, see Figure 6. Since we have less than $(\kappa_2/N)^{\frac{1}{d-2}}$ choices for $x_3, \ldots, x_d$, at most N choices for x_1 and the value of x_2 is adjusted to satisfy the constraint $\sum_{i=1}^d x_i = N$, the size of B_N is less than κ_2. Hence, B_N cannot contain all vertices of Y_N. Moreover, no vertices of Y_N lie on $x_1 x_2$-plane, so there exists some vertex $u \in Y_N$ such that $u_j \geq (\kappa_2/N)^{\frac{1}{d-2}} - 1$ for some $j \in [3..d]$, which is in $T(v) \cap L_N$ for some $v \in V$. ◀

The existence of this vertex v is the root of the problem. We conclude with the following statement.

▶ **Lemma 19.** *Any CDR defined on $\mathcal{G}_N^+ \subset \mathbb{Z}^d$ with κ_2 inner leaves in $T_{x_1 x_2}$ between $L_{N/2}$ and L_N has $\Omega((\kappa_2/N)^{\frac{1}{d-2}})$ error.*

Proof. Apply Lemma 18 to obtain a vertex $v \in L_{N/2-1} \cap T_{x_1 x_2}$ that satisfies some $u \in T(v) \cap L_N$ with $u_j \geq (\kappa_2/N)^{\frac{1}{d-2}} - 1$ for some $j \in [3..d]$. Let u' be the intersection of $\overline{ou}$ and the affine plane containing $L_{N/2-1}$, see Figure 6. As L_N and $L_{N/2-1}$ are parallel, $\frac{u'_j - o_j}{u_j - o_j} \geq \frac{1}{3}$ for $N \geq 6$, this implies that $u'_j = \Omega((\kappa_2/N)^{\frac{1}{d-2}})$. By construction, we have that v is on the $dig(o,u)$ and $v_j = 0$, hence $|| \cdot ||_\infty$ distance between $dig(o,u)$ and $\overline{ou}$ is $\Omega((\kappa_2/N)^{\frac{1}{d-2}})$. ◄

Combining with Theorem 3 gives us a lower bound for CDRs in d dimensions.

▶ **Theorem 4.** *Any CDR in $\mathbb{Z}^d$ has $\Omega(\log^{1/(d-1)} N)$ error.*

Proof. By Theorem 3 and Lemma 19, the error is $\Omega(\frac{N \log N}{N + \kappa_2})$ and $\Omega((\kappa_2/N)^{\frac{1}{d-2}})$, where κ_2 is the number of inner leaves in $T_{x_1 x_2}$ between $L_{N/2}$ and L_N. The balance between the two is obtained by choosing $\kappa_2 = \Theta(N \log^{\frac{d-2}{d-1}} N)$, giving the $\Omega(\log^{1/(d-1)} N)$ lower bound. ◄

6 Final remarks

Common intuition would say that the $\Omega(\log N)$ lower bound for the error of two-dimensional CDR and CDS automatically extends to higher dimensions. The observation that this is not true opens up new ways in which research can continue. We believe that further analysis of the mapping between the three spaces (from CDR in high dimensions to the 2-D weak CDR to the two-colored pointset) and the high interdependence between the three spaces can help in designing better lower and upper bounds.

Our lower bound $\Omega(\log^{1/(d-1)} N)$ extends the previous lower bound. The next step would be to close the gap between $\Omega(\log^{1/2} N)$ and $O(\log N)$ bounds in three dimensions. Even if the final answer ends up being $\Theta(\log N)$ we believe that the relationship between high dimensional CDRs, weak CDRs induced in subspaces and the mapping to pointset gives a better understanding of CDRs.

We also find that weak CDRs are an interesting research topic on their own. In particular, we would like to find the relationship between the number of inner leaves and the error of the construction. That is, say that we want a CDR with $O(e)$ error (for some $e \leq \log n$). What is the minimum number of leaves $\ell = \ell(e)$ that such a CDR must have? Can we find such a construction?

Theorem 3 seems to indicate a linear relationship between the two, and it is not hard to obtain one (an example is given in [3]). However, this construction is most likely not the best possible one. Indeed, even if we are interested in $O(\log N)$ error, this construction creates a large number of inner leaves, but we know of CDRs with the same error and no inner leaves. Thus, the question becomes, can we significantly improve upon the greedy construction, which can be found in the full version of the paper in [3]? Or is there some exponential dependency between the number of inner leaves and the error of the weak CDR?

───── **References** ─────

1   Tetsuo Asano, Danny Z. Chen, Naoki Katoh, and Takeshi Tokuyama. Efficient algorithms for optimization-based image segmentation. *International Journal of Computational Geometry and Applications*, 11(2):145–166, 2001.
2   Man-Kwun Chiu and Matias Korman. High dimensional consistent digital segments. *SIAM Journal on Discrete Mathematics*, 32(4):2566–2590, 2018.
3   Man-Kwun Chiu, Matias Korman, Martin Suderland, and Takeshi Tokuyama. Distance bounds for high dimensional consistent digital rays and 2-d partially-consistent digital rays, 2020. `arXiv:2006.14059`.

4   Iffat Chowdhury and Matt Gibson. A characterization of consistent digital line segments in $\mathbb{Z}^2$. In Nikhil Bansal and Irene Finocchi, editors, *Proceedings of the 23rd Annual European Symposium on Algorithms*, volume 9294, pages 337–348, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg.

5   Iffat Chowdhury and Matt Gibson. Constructing consistent digital line segments. In Evangelos Kranakis, Gonzalo Navarro, and Edgar Chávez, editors, *Proceedings of the 12th Latin American Theoretical Informatics Symposium*, volume 9644, pages 263–274, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.

6   Tobias Christ, Dömötör Pálvölgyi, and Miloš Stojaković. Consistent digital line segments. *Discrete & Computational Geometry*, 47(4):691–710, 2012.

7   Jinhee Chun, Natsuda Kaothanthong, Ryosei Kasai, Matias Korman, Martin Nöllenburg, and Takeshi Tokuyama. Algorithms for computing the maximum weight region decomposable into elementary shapes. *Computer Vision and Image Understanding*, 116(7):803–814, 2012.

8   Jinhee Chun, Matias Korman, Martin Nöllenburg, and Takeshi Tokuyama. Consistent digital rays. *Discrete and Computational Geometry*, 42(3):359–378, 2009.

9   Jacob E. Goodman, Richard Pollack, and Bernd Sturmfels. Coordinate representation of order types requires exponential storage. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 405–410. ACM, 1989.

10   Reinhard Klette and Azriel Rosenfeld. Digital straightness—a review. *Discrete Appl. Math.*, 139(1–3):197–230, 2004.

11   Michael G. Luby. Grid geometries which preserve properties of Euclidean geometry: A study of graphics line drawing algorithms. In *NATO Conference on Graphics/CAD*, pages 397–432, 1987.

12   Jiří Matoušek. *Geometric Discrepancy: An Illustrated Guide*. Algorithms and Combinatorics. Springer Berlin Heidelberg, 1999. URL: `https://books.google.co.jp/books?id=BKvXj1GisP0C`.

13   Wolfgang Schmidt. Irregularities of distribution, vii. *Acta Arithmetica*, 21(1):45–50, 1972. URL: `http://eudml.org/doc/205130`.